

<b>REPORT DOCUMENTATION PAGE</b>			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 24-09-2018		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 22-Sep-2009 - 30-Sep-2015	
4. TITLE AND SUBTITLE Final Report: A Cyber Awareness Framework for Attack Analysis, Prediction, and Visualization			5a. CONTRACT NUMBER W911NF-09-1-0553		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611103		
6. AUTHORS			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of California - Santa Barbara 3227 Cheadle Hall 3rd floor, MC 2050 Santa Barbara, CA 93106 -2050			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 56142-CS-MUR.126		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Richard Kemmerer
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 805-893-4232

# RPPR Final Report

as of 04-Dec-2018

Agency Code:

Proposal Number: 56142CSMUR

Agreement Number: W911NF-09-1-0553

## INVESTIGATOR(S):

**Name:** Ph.D. Richard A. Kemmerer kemm@cs.uc

**Email:** kemm@cs.ucsb.edu

**Phone Number:** 80589342320000

**Principal:** Y

Organization: **University of California - Santa Barbara**

Address: 3227 Cheadle Hall, Santa Barbara, CA 931062050

Country: USA

DUNS Number: 094878394

EIN: 956006145W

**Report Date:** 31-Dec-2015

Date Received: 24-Sep-2018

**Final Report** for Period Beginning 22-Sep-2009 and Ending 30-Sep-2015

**Title:** A Cyber Awareness Framework for Attack Analysis, Prediction, and Visualization

**Begin Performance Period:** 22-Sep-2009

**End Performance Period:** 30-Sep-2015

**Report Term:** 0-Other

Submitted By: Ph.D. Richard Kemmerer

Email: kemm@cs.ucsb.edu

Phone: (805) 893-42320000

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

**STEM Degrees:** 5

**STEM Participants:** 7

**Major Goals:** The objective of this research is to develop novel situation awareness theories and techniques to obtain an accurate view of the available cyber-assets and to automatically determine the assets required to carry out each mission task. Based on this information, we will automatically assess the damage of attacks, possible next moves, and the impact on the missions. We will also model the behavior of adversaries to predict the threat of future attacks to the success of a mission. Finally, we will present the status of the current missions and the impact of possible countermeasures to a security officer, using a semantically-rich environment. Each of these technologies will be integrated into a coherent cyber-situation awareness framework.

Our approach is based on the following five main thrusts:

THRUST 1: Obtaining an up-to-date view of the available cyber-assets

=====

It is impossible to understand the impact of an attack without knowing what cyber-assets were compromised and what role in the mission they play. Therefore, an important prerequisite for situation awareness is the knowledge of the cyber-infrastructure assets on which the mission relies.

THRUST 2: Obtaining dependencies between missions and assets

=====

Cyber-missions are sets of tasks that must be performed in a specified order and within a specified time frame to enable and to support operational missions. These tasks can assume a hierarchical relationship. That is, there are high-level tasks that are composed of a series of more specific, low-level tasks, which could possibly be further broken down into even lower-level ones.

THRUST 3: Obtaining an accurate view of the impact of cyber-attacks

=====

Based on the available information on the cyber-assets that are required by each mission, our framework will draw meaningful conclusions about the current status of the missions being executed and the threats that different attacks pose. This analysis should generate a number of possible courses of action (COAs), highlighting cost-benefit trade offs.

THRUST 4: Obtaining actionable cyber-attack forecasts

=====

The cyber-awareness domain possesses significant challenges to mainstream game theory, which can only be

## RPPR Final Report as of 04-Dec-2018

overcome through fundamental research in this area. While it will often be possible to determine unequivocally whether or not a particular task was successfully accomplished, this may not always be the case. Conversely, the adversary may also not be able to determine accurately the current state of the mission. In fact, the mission's success may to a great extent rely on this. Partial information games, in which one or both players are not fully aware of the current "state" of the game, are especially challenging and, while optimal solutions may be impractical, it is possible to construct solutions that are suitable for cyber-attack forecasting. Another important challenge is that the complete set of adversary actions (i.e., attacks) will typically not be known a-priori. Additionally, one may be faced with large uncertainty regarding the information available to the adversaries and even the ultimate intent behind their actions.

### THRUST 5: Obtaining a semantically-rich, easy-to-grasp view of the cybermission status

=====

Successful information transfer to a decision-maker and the right analysis tools given a particular user-context are crucial components of superior cybersecurity situational awareness. Collecting, extracting, mapping, filtering, modeling, and predicting security data are all important steps for creating a safe, secure, and resilient cyberinfrastructure to support and protect important missions. But without the consideration of context-specific knowledge dissemination theory and mechanisms, the decision maker will not be able to arrive at the correct analysis, understanding, and dynamic explanation of such data, or unveil the options and consequences for the decisions ahead, and do all of this in a timely fashion before the window of opportunity for appropriate counter-measures closes. An effective approach to true cybersecurity-awareness requires a multi-disciplinary collaboration among security researchers and practitioners, game theoreticians, and visualization and user interface experts. There is a need and distinct opportunity for scientific advances in interactive visualization methodology that scales appropriately with regard to changes in display and interaction capabilities, and mission context.

**Accomplishments:** We present a high-level summary of our accomplishments over the final reporting period, organized by research thrust. For more information, see our extended report, appended in the uploads section.

### THRUST 1: Obtaining an up-to-date view of the available cyber-assets

=====

Our efforts have focused on developing deeper capabilities to understand the activity of assets that operate solely within the confines of enterprise networks, a domain that has seen little work by the broader community due to the difficulties of obtaining the necessary access to such traffic. In addition, we have undertaken development of an approach to accurately track assets even in the presence of aliasing of their network-visible identities.

### THRUST 2: Obtaining dependencies between missions and assets

=====

We continued tuning the Sarsia system to better detect dependencies. We also refined and tuned Rippler, which is a complementary system that extracts dependencies based on timing patterns. We continued our research on incorporating external information to augment local network perspectives, and have nearly completed our system infrastructure for very large scale archiving of enterprise activity for later use, such as for "what-if" analysis. We also undertook the development of a novel framework for computing statistical properties of network traffic using distributed sensing.

### THRUST 3: Obtaining an accurate view of the impact of cyber-attacks

=====

Our novel framework for assessing the impact of cyber-attacks handles triaging, and has been tested on a realistic, phase-based cyber-warfare exercise. We developed a system, dubbed Nazca, for picturing the download-behavior of executables within large-scale networks. We presented another system, Hulk, to detect malicious behaviors within browser extensions. Finally, we interviewed and began testing real-world industrial control networks to glean more information on realistic attack-preparedness and awareness.

### THRUST 4: Obtaining actionable cyber-attack forecasts

=====

We extended and refined our approaches to asymmetric games, addressed the solution of large-scale games using randomized methods, proposed some optimization-based algorithms for attack prediction, developed some measurement models for estimating the status of cyber assets, and developed some new approaches to the

## **RPPR Final Report**

as of 04-Dec-2018

consensus problem when under attack.

THRUST 5: Obtaining a semantically-rich, easy-to-grasp view of the cybermission status

=====

Over the past review period, we have made significant progress in three overall areas: 1) We updated our surround-view immersive situation room and visualization chamber, the UCSB Allosphere, to full-surround projection of 2D, 3D, and multimodal information displays. 2) We have created flexible tools for visual and multimodal cybersecurity data analysis. 3) We have been working towards automatic generation of user interfaces for different display platforms by conceptualizing a universal syntax for data and user interfaces, providing structured editing tools for data and UIs, and targeting transparent "deconstructible" user interfaces and visualizations.

**Training Opportunities:** Nothing to Report

## RPPR Final Report as of 04-Dec-2018

**Results Dissemination:** Press/Media information about the SecLab's success at the DARPA Cyber Grand Challenge can be found here:

-<https://www.darpa.mil/news-events/2015-07-08>.

-<http://archive.darpa.mil/cybergrandchallenge/>

-<http://www.news.ucsb.edu/2015/015786/ultimate-capture-flag-game>

A Norwegian article highlighting our research on cyber-security and the power industry - September 2, 2014. <https://frifagbevegelse.no/article-6.158.167601.7bc679380b>

Conferences and meetings during reporting period where results were presented:

Dhilung Kirat, Giovanni Vigna, Christopher Kruegel. (2014 August). BareCloud: Bare-metal Analysis-based Evasive Malware Detection. 23rd USENIX Security Symposium, San Diego, CA, USA.

Alexandros Kapravelos, Chris Grier, Neha Chachra, Christopher Kruegel, Giovanni Vigna, Vern Paxson (2014, August). Hulk: Eliciting Malicious Behavior in Browser Extensions. 23rd USENIX Security Symposium, San Diego, CA, USA.

Yinzhi Cao, Yan Shoshitaishvili, Kevin Borgolte, Christopher Kruegel, Giovanni Vigna, Yan Chen. (2014, September). Protecting Web-based Single Sign-on Protocols against Relying Party Impersonation Attacks through a Dedicated Bi-directional Authenticated Secure Channel. International Symposium on Research in Attacks, Intrusions and Defenses (RAID). Gothenburg, Sweden.

Jacopo Corbetta, Luca Invernizzi, Christopher Kruegel, Giovanni Vigna. (2014, September). Eyes of a Human, Eyes of a Program: Leveraging different views of the web for analysis and detection. International Symposium on Research in Attacks, Intrusions and Defenses (RAID). Gothenburg, Sweden.

Giovanni Vigna. (2014, October). Keynote address at the IEEE Malware Conference in Puerto Rico. Fajardo, Puerto Rico.

Giovanni Vigna. (2014, November). Distinguished lecture on research related to cyber security and situation awareness. University of Illinois, Chicago, IL, USA.

Donghao Ren, Tobias Höllerer, Xiaoru Yuan (2014, November). iVisDesigner: Expressive Interactive Design of Information Visualizations. VIS 2014. Paris, France.

Charles Roberts, Matthew Wright, JoAnn Kuchera-Morin, Tobias Höllerer. (2014, November). Gibber: Abstractions for Creative Multimedia Programming ACM Multimedia 2014. Orlando, FL, USA.

Kyriakos Vamvoudakis. (2014, December). Presentation at University of Southern California, related to research findings in the Cybaware MURI. Los Angeles, CA, USA.

David A Copp, Joao Hespanha. (2014, December). Nonlinear Output-Feedback Model Predictive Control with Moving Horizon Estimation. 53rd IEEE Conference on Decision and Control. Los Angeles, CA, USA

Yinzhi Cao, Yanick Fratantonio, Antonio Bianchi, Manuel Egele, Christopher Kruegel, Giovanni Vigna, and Yan Chen. (2015, February). EdgeMiner: Automatically Detecting Implicit Control Flow Transitions through the Android Framework 2015 Network and Distributed System Security (NDSS) Symposium. San Diego, CA, USA.

Yan Shoshitaishvili, Fish Wang, Christophe Hauser, Christopher Kruegel, and Giovanni Vigna. (2015, February). Firmallice – Automatic Detection of Authentication Bypass Vulnerabilities in Binary Firmware. 2015 Network and Distributed System Security (NDSS) Symposium. San Diego, CA, USA.

Alexandros Kapravelos. (2015, February). Messaging Malware Mobile Anti-Abuse Working Group (M3AAWG). Hulk: Eliciting Malicious Behavior in Browser Extensions. San Francisco, CA, USA. February 15, 2015.

Antonio Bianchi, Jacopo Corbetta, Luca Invernizzi, Yanick Fratantonio, Christopher Kruegel, Giovanni Vigna. (2015, May). What the App is That? Deception and Countermeasures in the Android User Interface. 36th IEEE

# RPPR Final Report

## as of 04-Dec-2018

Symposium on Security and Privacy. San Jose, CA, USA.

Yanick Fratantonio, Aravind Machiry, Antonio Bianchi, Christopher Kruegel, and Giovanni Vigna. (2015, August). CLAPP: Characterizing Loops in Android Applications. 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering. Bergamo, Italy.

**Honors and Awards:** July 2015 -- The UCSB Hacking team, led by Co-PIs Giovanni Vigna and Christopher Kruegel, finished 3rd place in the DARPA Cyber Grand Challenge in Las Vegas, NV.

### Protocol Activity Status:

**Technology Transfer:** SecLab researchers coordinated and cooperated with DARPA agencies via the Cyber Grand Challenge in Summer 2015.

### PARTICIPANTS:

**Participant Type:** Graduate Student (research assistant)

**Participant:** Patrick Baxter

**Person Months Worked:** 3.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Enrico Bazzoli

**Person Months Worked:** 2.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Graduate Student (research assistant)

**Participant:** Jacopo Corbetta

**Person Months Worked:** 2.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Graduate Student (research assistant)

**Participant:** Luca Invernizzi

**Person Months Worked:** 4.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Tobias Jarmuzek

**RPPR Final Report**  
as of 04-Dec-2018

**Person Months Worked:** 4.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Pekko Lipsanen

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Jiaqui Liu

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Vasileios Mavroudis

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Marius Muench

**Person Months Worked:** 2.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Graduate Student (research assistant)

**Participant:** Ali Zand

**Person Months Worked:** 5.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Graduate Student (research assistant)

**RPPR Final Report**  
as of 04-Dec-2018

**Participant:** Christopher Hall  
**Person Months Worked:** 7.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**

**Participant Type:** Graduate Student (research assistant)

**Participant:** David Copp  
**Person Months Worked:** 2.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**

**Participant Type:** Graduate Student (research assistant)

**Participant:** Norma Saiph Savage  
**Person Months Worked:** 3.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**

**Participant Type:** Co PD/PI

**Participant:** Joao Hespanha  
**Person Months Worked:** 2.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**

**Participant Type:** Staff Scientist (doctoral level)

**Participant:** Kyriakos Vamvoudakis  
**Person Months Worked:** 10.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**

**Participant Type:** PD/PI

**Participant:** Richard A Kemmerer  
**Person Months Worked:** 4.00  
Project Contribution:  
International Collaboration:  
International Travel:  
National Academy Member: N  
Other Collaborators:

**Funding Support:**



# RPPR Final Report

as of 04-Dec-2018

**Participant Type:** Non-Student Research Assistant

**Participant:** Timothy Robinson

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

**Participant Type:** Non-Student Research Assistant

**Participant:** Nikolaos Melissaris

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

## ARTICLES:

**Publication Type:** Journal Article

Peer Reviewed: Y

**Publication Status:** 1-Published

**Journal:** Journal of Computer Security

Publication Identifier Type:

Publication Identifier:

Volume: 18

Issue: 5

First Page #: 861

Date Submitted:

Date Published:

Publication Location:

**Article Title:** Static Analysis for Detecting Taint-style Vulnerabilities in Web Applications

**Authors:**

**Keywords:** Program analysis, static analysis, data flow analysis, alias analysis, web application security, scripting languages security, cross-site scripting, SQL injection, PHP

**Abstract:** The number and the importance of web applications have increased rapidly over the last years. At the same time, the quantity and impact of security vulnerabilities in such applications have grown as well. Since manual code reviews are time-consuming, error-prone and costly, the need for automated solutions has become evident. In this paper, we address the problem of vulnerable web applications by means of static source code analysis. More precisely, we use flow-sensitive, interprocedural and context-sensitive data flow analysis to discover vulnerable points in a program. In addition to the taint analysis at the core of our engine, we employ a precise alias analysis targeted at the unique reference semantics commonly found in scripting languages. Moreover, we enhance the quality and quantity of the generated vulnerability reports by employing an iterative two-phase algorithm for fast and precise resolution of file inclusions. The presented concepts are targeted at the general class of

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** IEEE Transactions on Visualization and Computer Graphics  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 16      **Issue:** 6      **First Page #:** 0  
**Date Submitted:**      **Date Published:**  
**Publication Location:**

**Article Title:** behaviorism: a Framework for Dynamic Data Visualization

**Authors:**

**Keywords:** Frameworks, information visualization, information art, dynamic data.

**Abstract:** While a number of information visualization software frameworks exist, creating new visualizations, especially those that involve novel visualization metaphors, interaction techniques, data analysis strategies, and specialized rendering algorithms, is still often a difficult process. To facilitate the creation of novel visualizations we present a new software framework, behaviorism, which provides a wide range of flexibility when working with dynamic information on visual, temporal, and ontological levels, but at the same time providing appropriate abstractions which allow developers to create prototypes quickly which can then easily be turned into robust systems. The core of the framework is a set of three interconnected graphs, each with associated operators: a scene graph for high-performance 3D rendering, a data graph for different layers of semantically-linked heterogeneous data, and a timing graph for sophisticated control of scheduling, interaction, and animation. In particular,

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** IEEE Security and Privacy Magazine  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 9      **Issue:** 1      **First Page #:** 64  
**Date Submitted:**      **Date Published:**  
**Publication Location:**

**Article Title:** Analysis of a Botnet Takeover

**Authors:**

**Keywords:** botnets, Torpig, hacking

**Abstract:** This article describes an effort to take control of a particularly sophisticated and insidious botnet and study its operations for a period of 10 days. It summarizes what the authors learned and reports on what has happened to that botnet since. Botnets, networks of malware-infected machines (bots) controlled by an adversary, are the root cause of a large number of Internet security problems. They're the primary way cybercriminals carry out their nefarious tasks, such as sending spam, launching denial-of-service attacks, or stealing personal data. A particularly sophisticated and insidious variety is called Torpig, malware designed to harvest sensitive information such as bank account and credit-card data from its victims.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** ACM Transactions on Intelligent Systems and Technology  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 0      **Issue:** 0      **First Page #:** 0  
**Date Submitted:**      **Date Published:**  
**Publication Location:**

**Article Title:** TopicNets: Visual Analysis of Large Text Corpora with Topic Modeling

**Authors:**

**Keywords:** I.2.7 [Text analysis]: Natural Language Processing; H.5.3 [Web-based interaction]: Group and Organization Interfaces, Topic Modeling, Text Visualization, Graph Visualization

**Abstract:** We present TopicNets, a web-based system for visual and interactive analysis of large sets of documents using statistical topic models. A range of visualization types and control mechanisms to support knowledge discovery are presented. These include corpus and document specific views, iterative topic modeling, search, and visual altering. Drill-down functionality is provided to allow analysts to visualize individual document sections and their relations within the global topic space. Analysts can search across a data set through a set of expansion techniques on selected document and topic nodes. Furthermore, analysts can select relevant subsets of documents and perform real-time topic modeling on these subsets to interactively visualize topics at various levels of granularity, allowing for a better understanding of the documents. A discussion of the design and implementation choices for each visual analysis technique is presented. This is followed by a discussion of three diverse use cases.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** IEEE Transactions on Information Forensics and Security  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 6      **Issue:** 1      **First Page #:** 175  
**Date Submitted:**      **Date Published:**  
**Publication Location:**

**Article Title:** Towards Situational Awareness of Large-scale Botnet Probing Events

**Authors:**

**Keywords:** Botnet, Global property extrapolation, Honeynet, Scan strategy inference, Situational awareness, Statistical inference

**Abstract:** Botnets dominate today's attack landscape. In this work we investigate ways to analyze collections of malicious probing traffic in order to understand the significance of large-scale "botnet probes". In such events, an entire collection of remote hosts together probes the address space monitored by a sensor in some sort of coordinated fashion. Our goal is to develop methodologies by which sites receiving such probes can infer -- using purely local observation -- information about the probing activity: What scanning strategies does the probing employ? Is this an attack that specifically targets the site, or is the site only incidentally probed as part of a larger, indiscriminant attack? Our analysis draws upon extensive honeynet data to explore the prevalence of different types of scanning, including properties such as trend, uniformity, coordination, and darknet avoidance. In addition, we design schemes to extrapolate the global properties of scanning events (e.g., total population and

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** IEEE Transactions on Visualization and Computer Graphics  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 0      **Issue:** 0      **First Page #:** 0  
**Date Submitted:**      **Date Published:**  
**Publication Location:**

**Article Title:** Stereoscopic Highlighting: 2D Graph Visualization on Stereo Displays

**Authors:**

**Keywords:** graph visualization, stereo displays, virtual reality.

**Abstract:** In this paper we present a new technique and prototype graph visualization system, stereoscopic highlighting, to help answer accessibility and adjacency queries when interacting with a node-link diagram. Our technique utilizes stereoscopic depth to highlight regions of interest in a 2D graph by projecting these parts onto a plane closer to the viewpoint of the user. This technique aims to isolate and magnify specific portions of the graph that need to be explored in detail without resorting to other highlighting techniques like color or motion, which can then be reserved to encode other data attributes. This mechanism of stereoscopic highlighting also enables focus+context views by juxtaposing a detailed image of a region of interest with the overall graph, which is visualized at a further depth with correspondingly less detail. In order to validate our technique, we ran a controlled experiment with 16 subjects comparing static visual highlighting to stereoscopic highlighting on 2D and

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** Control of Cyber-Physical Systems, Lecture Notes in Control and Information Sciences  
**Publication Identifier Type:**      **Publication Identifier:**  
**Volume:** 449      **Issue:** 0      **First Page #:** 85  
**Date Submitted:**      **Date Published:**  
**Publication Location:** Berlin

**Article Title:** Formulating Cyber-Security as Convex Optimization Problems

**Authors:**

**Keywords:** Cyber-Security, Convex Optimization, System Identification, iCTF

**Abstract:** Mission-centric cyber-security analysts require a complete overview and understanding of the state of a mission and any potential threats to their completion. To facilitate this, we propose optimization-based algorithms that can be used to predict in real-time how an attacker may try to compromise a cyber-mission with a limited amount of resources, based on a model that takes into account potential damage to the mission and probabilistic uncertainty. Two different optimization schemes are considered: one where all the mission data is known a priori to the attacker and another where system identification and a moving horizon optimization is used to produce the estimates based on historical data. Our schemes are compared with real attacks carried out by human players in the 2011 international Capture The Flag (iCTF) hacking competition.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Control of Cyber-Physical Systems, Lecture Notes in Control and Information Sciences

Publication Identifier Type:      Publication Identifier:

Volume: 449      Issue: 0      First Page #: 65

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Cyber-attack Forecast Modeling and Complexity Reduction Using a Game Theoretic Framework

**Authors:**

**Keywords:** cyber-attack forecasting, cyber-security

**Abstract:** The security community has placed a significant emphasis on developing tools and techniques to address known security issues. Some examples of this emphasis include security tools such as anti-virus software and Intrusion Detection Systems (IDS). This reactive approach to security is effective against novice adversaries (i.e. script kiddies) because they typically use off-the-shelf tools and popular techniques to conduct their attacks. In contrast, the innovative adversaries often devise novel attack vectors and methodologies that can render reactive measures inadequate. These pioneering adversaries have continually pushed the security frontier forward and motivate a need for proactive security approaches. A proactive approach that we pursue in this research is actionable cyber-attack forecasting. The objectives of actionable cyber-attack forecasting are to learn an attacker's behavioral model, to predict future attacks, and to select appropriate countermeasures. The computational comp

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Automatica

Publication Identifier Type:      Publication Identifier:

Volume: 49      Issue: 5      First Page #: 0

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Randomized Sampling for Large Zero-Sum Games

**Authors:**

**Keywords:** Game theory, Randomized algorithms, Zero-Sum Games, Optimization

**Abstract:** This paper addresses the solution of large zero-sum matrix games using randomized methods. We formalize a procedure, termed as the sampled security policy (SSP) algorithm, by which a player can compute policies that, with a high confidence, are security policies against an adversary using randomized methods to explore the possible outcomes of the game. The SSP algorithm essentially consists of solving a stochastically sampled subgame that is much smaller than the original game. We also propose a randomized algorithm, termed as the sampled security value (SSV) algorithm, which computes a high -confidence security-level (i.e., worst-case outcome) for a given policy, which may or may not have been obtained using the SSP algorithm. For both the SSP and the SSV algorithms we provide results to determine how many samples are needed to guarantee a desired level of confidence. We start by providing results when the two players sample policies with the same distribution and subsequently extend

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: N      **Publication Status:** 5-Submitted

**Journal:** Proceedings of the ACM Conference on Computer and Communications Security

Publication Identifier Type:      Publication Identifier:

Volume: 0      Issue: 0      First Page #: 0

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Shady Paths: Leveraging Surfing Crowds to Detect Malicious Web Pages

**Authors:**

**Keywords:** Security, HTTP redirections, detection, malware

**Abstract:** The web is one of the most popular vectors to spread malware. Attackers lure victims to visit compromised web pages or entice them to click on malicious links. These victims are redirected to sites that exploit their browsers or trick them into installing malicious software using social engineering. In this paper, we tackle the problem of detecting malicious web pages from a novel angle. Instead of looking at particular features of a (malicious) web page, we analyze how a large and diverse set of web browsers reach these pages. That is, we use the browsers of a collection of web users to record their interactions with websites, as well as the redirections they go through to reach their final destinations. We then aggregate the different redirection chains that lead to a specific web page and analyze the characteristics of the resulting redirection graph. As we will show, these characteristics can be used to detect malicious pages. We argue that our approach is less prone to evasion than

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: N      **Publication Status:** 5-Submitted

**Journal:** Proceedings of the ACM Conference on Computer and Communications Security (CCS 2013)

Publication Identifier Type:      Publication Identifier:

Volume: 0      Issue: 0      First Page #: 0

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Detecting Stealthy, Distributed SSH Brute-Forcing

**Authors:**

**Keywords:** Scanning; SSH; Brute-forcing; Distribute

**Abstract:** In this work we propose a general approach for detecting distributed malicious activity in which individual attack sources each operate in a stealthy, low-profile manner. We base our approach on observing statistically significant changes in a parameter that summarizes aggregate activity, bracketing a distributed attack in time, and then determining which sources present during that interval appear to have coordinated their activity. We apply this approach to the problem of detecting stealthy distributed SSH bruteforcing activity, showing that we can model the process of legitimate users failing to authenticate using a beta-binomial distribution, which enables us to tune a detector that trades off an expected level of false positives versus time-to-detection. Using the detector we study the prevalence of distributed bruteforcing, finding dozens of instances in an extensive 8-year dataset collected from a site with several thousand SSH users. Many of the attacks—some of which last month

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: N      **Publication Status:** 5-Submitted

**Journal:** Proc. 52nd IEEE Conference on Decision and Control

Publication Identifier Type:      Publication Identifier:

Volume: 0      Issue: 0      First Page #: 0

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Gossip Average Consensus in a Byzantine Environment Using Stochastic Set-Valued Observers

**Authors:**

**Keywords:** Byzantine faults

**Abstract:** We address the problem of a consensus system in the presence of Byzantine faults seen as an attacker injecting a perturbation in the state of the nodes. We propose the use of Set-Valued Observers to detect if the state observations are compatible with the system dynamics. The method is extended to the stochastic case by introducing a strategy to construct a set that is guaranteed to contain all possible states with, at least, a pre-specified desired probability. The proposed algorithm is stable in the sense that it requires a finite number of vertices to represent polytopic sets while also enabling the a priori computation of the largest magnitude of a disturbance that an attacker can inject without being detected

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** IEEE Computer Graphics & Applications

Publication Identifier Type:      Publication Identifier:

Volume: 33      Issue: 6      First Page #: 14

Date Submitted:      Date Published:

Publication Location:

**Article Title:** Spatial Interaction in a Multiuser Immersive Instrument

**Authors:**

**Keywords:** immersive environments; HCI; virtual environment

**Abstract:** The appropriate support of spatial interaction is a perennial challenge in all kinds of VR environments. However, the results can be especially rewarding when you're interacting alongside other users in a surround-view and surround-sound immersive environment, such as the AlloSphere at the University of California, Santa Barbara ([www.allosphere.ucsb.edu](http://www.allosphere.ucsb.edu)). The AlloSphere, conceived by JoAnn Kuchera-Morin, is a large scientific and artistic instrument for immersive human-centered visualization, sonification (using nonspeech audio to present information), and interactive data manipulation. We allude to both the scientific and musical senses of "instrument." Like a microscope or telescope, the AlloSphere makes new realms accessible to human perception. Like the musical instruments in an orchestra or ensemble, it aims to facilitate multiuser parametric control of complex information.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Computers & Graphics

Publication Identifier Type:

Publication Identifier:

Volume: 40

Issue: 1

First Page #: 10

Date Submitted:

Date Published:

Publication Location:

**Article Title:** Immersive Full-Surround Multi-User System Design

**Authors:**

**Keywords:** VR systems; Display technology; Multi-user; Multimodal interaction; Immersion

**Abstract:** This paper describes our research in full-surround, multimodal, multi-user, immersive instrument design in a large VR instrument. The three-story instrument, designed for large-scale, multimodal representation of complex and potentially high-dimensional information, specifically focuses on multi-user participation by facilitating interdisciplinary teams of co-located researchers in exploring complex information through interactive visual and aural displays in a full-surround, immersive environment. We recently achieved several milestones in the instrument's design that improves multi-user participation when exploring complex data representations and scientific simulations. These milestones include affordances for "ensemble-style" interaction allowing groups of participants to see, hear, and explore data as a team using our multi-user tracking and interaction systems; separate visual display modes for rectangular legacy content and for seamless surround-view stereoscopic projection usin

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support:

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** IEEE Transactions on Automatic Control

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TAC.2014.2351671

Volume: 59

Issue: 12

First Page #: 3209

Date Submitted: 9/21/18 12:00AM

Date Published: 12/1/14 8:00AM

Publication Location:

**Article Title:** Detection in Adversarial Environments

**Authors:** Kyriakos G. Vamvoudakis, Joao P. Hespanha, Bruno Sinopoli, Yilin Mo

**Keywords:** Adversarial detection, byzantine sensors, cyber security, zero-sum games, estimation.

**Abstract:** We propose new game theoretic approaches to estimate a binary random variable based on sensor measurements that may have been corrupted by a cyber-attacker. The estimation problem is formulated as a zero-sum partial information game in which a detector attempts to minimize the probability of an estimation error and an attacker attempts to maximize this probability. While this problem can be solved exactly by reducing it to the computation of the value of a matrix, this approach is computationally feasible only for a small number of sensors. The two key results of this paper provide complementary computationally efficient solutions to the construction of the optimal detector. The first result provides an explicit formula for the optimal detector but it is only valid when the number of sensors is roughly smaller than two over the probability of sensor errors. In contrast, the detector provided by the second result is valid for an arbitrary number of sensor. While it may ...

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y



## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** IEEE Transactions on Signal Processing

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TSP.2013.2284145

Volume: 62

Issue: 1

First Page #: 31

Date Submitted: 9/21/18 12:00AM

Date Published: 1/1/14 8:00AM

Publication Location:

**Article Title:** Resilient Detection in the Presence of Integrity Attacks

**Authors:** Yilin Mo, Joao P. Hespanha, Bruno Sinopoli

**Keywords:** Detection algorithms, robustness, fault tolerance

**Abstract:** We consider the detection of a binary random state based on measurements that can be manipulated by an attacker. The attacker is assumed to have full information about the true value of the state to be estimated as well as the values of all the measurements. However, the attacker can only manipulate of the measurements. The detection problem is formulated as a minimax optimization, where one seeks to construct an optimal detector that minimizes the “worst-case” probability of error against all possible manipulations by the attacker. We show that if the attacker can manipulate at least half the measurements then the optimal worst-case detector should ignore all measurements and be based solely on the a-priori information. When the attacker can manipulate less than half of the measurements, we show that the optimal detector is a threshold rule based on a Hamming-like distance between the (manipulated) measurement vector and two appropriately defined sets....

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** IEEE Transactions on Visualization and Computer Graphics

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TVCG.2014.2346291

Volume: 20

Issue: 12

First Page #: 2092

Date Submitted: 9/21/18 12:00AM

Date Published: 12/1/14 8:00AM

Publication Location:

**Article Title:** iVisDesigner: Expressive Interactive Design of Information Visualizations

**Authors:** Donghao Ren, Tobias Hollerer, Xiaoru Yuan

**Keywords:** Visualization design, Interactive Design, Interaction, Expressiveness, Web-based visualization

**Abstract:** We present the design, implementation and evaluation of iVisDesigner, a web-based system that enables users to design information visualizations for complex datasets interactively, without the need for textual programming. Our system achieves high interactive expressiveness through conceptual modularity, covering a broad information visualization design space. iVisDesigner supports the interactive design of interactive visualizations, such as provisioning for responsive graph layouts and different types of brushing and linking interactions. We present the system design and implementation, exemplify it through a variety of illustrative visualization designs and discuss its limitations. A performance analysis and an informal user study are presented to evaluate the system.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

### CONFERENCE PAPERS:

**Publication Type:** Conference Paper or Presentation      **Publication Status:** 1-Published

**Conference Name:** USENIX Security

Date Received: 21-Sep-2018

Conference Date: 20-Aug-2014

Date Published: 21-Aug-2014

Conference Location: San Deigo, CA

**Paper Title:** Hulk: Eliciting Malicious Behavior in Browser Extensions

**Authors:** Alex Kapravelos, Chris Grier, Neha Chachra, Christopher Kruegel, Giovanni Vigna, Vern Paxson

Acknowledged Federal Support: Y

**RPPR Final Report**  
as of 04-Dec-2018

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** USENIX Security  
Date Received: 21-Sep-2018 Conference Date: 21-Aug-2014 Date Published: 20-Aug-2014  
Conference Location: San Diego, CA  
**Paper Title:** Ten Years of iCTF: The Good, The Bad, and The Ugly  
**Authors:** Giovanni Vigna, Kevin Borgolte, Jacopo Corbetta, Adam Doupé, Yanick Fratantonio, Luca Invernizzi, Dr  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** USENIX Security  
Date Received: 21-Sep-2018 Conference Date: 20-Aug-2014 Date Published: 20-Aug-2014  
Conference Location: San Diego, CA  
**Paper Title:** BareCloud: Bare-metal Analysis-based Evasive Malware Detection  
**Authors:** Dhilung Kirat, Giovanni Vigna, Christopher Kruegel  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** RAID Symposium  
Date Received: 21-Sep-2018 Conference Date: 17-Sep-2014 Date Published: 17-Sep-2014  
Conference Location: Gothenburg, Sweden  
**Paper Title:** Eyes of a Human, Eyes of a Program: Leveraging Different Views of the Web for Analysis and Detection  
**Authors:** Jacopo Corbetta, Luca Invernizzi, Christopher Kruegel, Giovanni Vigna  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** RAID Symposium  
Date Received: 21-Sep-2018 Conference Date: 17-Sep-2014 Date Published: 17-Sep-2014  
Conference Location: Gothenburg, Sweden  
**Paper Title:** Protecting Web Single Sign-on against Relying Party Impersonation Attacks through a Bi-directional Secure Channel with Authentication  
**Authors:** Yinzhì Cao, Yan Shoshitaishvili, Kevin Borgolte, Christopher Kruegel, Giovanni Vigna, Yan Chen  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** the ACM International Conference  
Date Received: 21-Sep-2018 Conference Date: 03-Nov-2014 Date Published:  
Conference Location: Orlando, Florida, USA  
**Paper Title:** Gibber: Abstractions for Creative Multimedia Programming  
**Authors:** Charles Roberts, Matthew Wright, JoAnn Kuchera-Morin, Tobias Höllerer  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** GLOBECOM 2014 - 2014 IEEE Global Communications Conference  
Date Received: 21-Sep-2018 Conference Date: 08-Dec-2014 Date Published:  
Conference Location: Austin, TX, USA  
**Paper Title:** Resilience of LTE networks against smart jamming attacks  
**Authors:** Farhan M. Azizy, Jeff S. Shamma, and Gordon L. Stuber  
Acknowledged Federal Support: **Y**

## RPPR Final Report as of 04-Dec-2018

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** 2014 IEEE 53rd Annual Conference on Decision and Control (CDC)  
Date Received: 21-Sep-2018 Conference Date: 15-Dec-2014 Date Published:  
Conference Location: Los Angeles, CA, USA  
**Paper Title:** LP formulation of asymmetric zero-sum stochastic games  
**Authors:** Lichun Li, Jeff Shamma  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** the 2014 ACM conference  
Date Received: 21-Sep-2018 Conference Date: 17-Aug-2014 Date Published:  
Conference Location: Chicago, Illinois, USA  
**Paper Title:** Native Actors: How to Scale Network Forensics  
**Authors:** Matthias Vallentin, Dominik Charousset, Thomas C. Schmidt, Vern Paxson, Matthias Wählisch  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** the 30th Annual Computer Security Applications Conference  
Date Received: 21-Sep-2018 Conference Date: 08-Dec-2014 Date Published:  
Conference Location: New Orleans, Louisiana  
**Paper Title:** Through the Eye of the PLC: Semantic Security Monitoring for Industrial Processes  
**Authors:** Dina Hadžiosmanovic, Robin Sommer, Emmanuele Zambon, Pieter H.Hartel  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** RAID Symposium  
Date Received: 21-Sep-2018 Conference Date: 17-Sep-2014 Date Published: 17-Sep-2014  
Conference Location: Gothenburg, Sweden  
**Paper Title:** Count Me In: Viable Distributed Summary Statistics for Securing High-Speed Networks  
**Authors:** Johanna Amann, Seth Hall, Robin Sommer  
Acknowledged Federal Support: **Y**

### DISSERTATIONS:

**Publication Type:** Thesis or Dissertation  
**Institution:**  
Date Received: 29-Aug-2012 Completion Date:  
**Title:** Defending Against Malicious Software  
**Authors:**  
Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation  
**Institution:**  
Date Received: 29-Aug-2012 Completion Date:  
**Title:** Interaction Methods for Large Scale Graph Visualization Systems – Using Manipulation to Aid Discovery  
**Authors:**  
Acknowledged Federal Support:

**RPPR Final Report**  
as of 04-Dec-2018

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2012

Completion Date:

**Title:** Where Are Malicious Networks Located?

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2012

Completion Date:

**Title:** A Mission-Centric Visualization Tool for Cybersecurity Situation Awareness

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 28-Aug-2013

Completion Date:

**Title:** CANDID: Classifying Assets in Networks by Determining Importance and Dependencies

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2013

Completion Date:

**Title:** Bridging Social and Semantic Computing -- Design and Evaluation of User Interfaces for Hybrid Systems

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2013

Completion Date:

**Title:** Bridging Dimensions in Visualization

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2013

Completion Date:

**Title:** Cybavis: an interactive tool to promote situational awareness in a mission centric cybersecurity scenario

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 26-Aug-2014

Completion Date:

**Title:** WMD: Non-Intrusive Host-Based Malware Detection

**Authors:**

Acknowledged Federal Support:

**RPPR Final Report**  
as of 04-Dec-2018

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 26-Aug-2014

Completion Date:

**Title:** Advanced Automated Web Application Vulnerability Analysis

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 26-Aug-2014

Completion Date:

**Title:** Stepping Up the Cybersecurity Game: Protecting Online Services from Malicious Activity

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 26-Aug-2014

Completion Date:

**Title:** Optimization in Stochastic Hybrid and Switching Systems

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2014

Completion Date:

**Title:** Towards Sound HTTP Request Causation Inference

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2014

Completion Date:

**Title:** Exploring Motion as a Modality for Visualizing Data

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2014

Completion Date:

**Title:** Immediacy In Creative Coding Environments

**Authors:**

Acknowledged Federal Support:

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2014

Completion Date:

**Title:** Emerging Methodologies for Interdisciplinary Research Practice

**Authors:**

Acknowledged Federal Support:

**RPPR Final Report**  
as of 04-Dec-2018

**Publication Type:** Thesis or Dissertation

**Institution:**

Date Received: 29-Aug-2014

Completion Date:

**Title:** Asymmetric Information Games and Cyber Security

**Authors:**

Acknowledged Federal Support:

# A Cyber Awareness Framework for Attack Analysis, Prediction, and Visualization

Funding Number: W911NF0910553

PI: Richard A. Kemmerer  
Co-PIs: Joao Hespanha  
Tobias Hollerer  
Christopher Kruegel  
Vern Paxson  
Jeff Shamma  
Giovanni Vigna

## Final Scientific Report

The objective of this research is to develop novel situation awareness theories and techniques to obtain an accurate view of the available cyber-assets and to automatically determine the assets required to carry out each mission task. Based on this information, we will automatically assess the damage of attacks, possible next moves, and the impact on the missions. We will also model the behavior of adversaries to predict the threat of future attacks to the success of a mission. Finally, we will present the status of the current missions and the impact of possible countermeasures to a security officer, using a semantically-rich environment. Each of these technologies will be integrated into a coherent cyber-situation awareness framework.

Our approach is based on the following five main thrusts:

- (1) an up-to-date view of the available cyber-assets
- (2) a comprehensive analysis of the dependencies between cyber-missions and cyber-assets
- (3) an accurate understanding of the impact of cyber-attacks
- (4) actionable cyber-attack forecasts
- (5) a semantically-rich, easy-to-grasp view of the cyber-mission status

We present our progress to date in each of these areas.

## THRUST 1: Obtaining an up-to-date view of the available cyber-assets

It is impossible to understand the impact of an attack without knowing what cyber-assets were compromised and what role in the mission they play. Therefore, an important prerequisite for situation awareness is the knowledge of the cyber-infrastructure assets on which the mission relies.

Our efforts have focused on developing deeper capabilities to understand the activity of assets that operate solely within the confines of enterprise networks, a domain that has seen little work by the broader community due to the difficulties of obtaining the necessary access to such traffic. In addition, we have undertaken development of an approach to accurately track assets even in the presence of aliasing of their network-visible identities.

### Enterprise Network Monitoring Capabilities

Gaining visibility into the assets and activity within an enterprise network requires grappling with the analysis of significantly different protocol than what is traditionally available, i.e. technology developed to monitor ingress/egress points between enterprises and the external Internet. To address this challenge, we developed support for monitoring several complex protocols that enterprises employ internally, providing a key capability to facilitate understanding of enterprise networks, including the difficulties that analysts will encounter when venturing into unknown organizations. Extending the open-source Bro network monitoring platform, we developed protocol support for SNMP and Radius, both crucial for understanding the semantics of internal communication. We have integrated these into Bro's recent public release. For the next release, we have begun to further integrate analyzers for SMB, MySQL, and Kerberos that were all originally developed externally by a 3rd party. SMB analysis in particular will provide unprecedented visibility into the activity of Windows systems, which today mostly escape such monitoring efforts.

In addition, we undertook developing visibility for the critical-infrastructure domain by pursuing one of the first studies of extracting and modelling application-level semantics from a network vantage point in real-world environments. Examining two weeks of activity from two operational water treatment plants serving a total of about one million people in two urban areas, we developed a network monitor that continuously tracks updates to process variables, as defined by their programmable logic controllers (PLCs). We then derived variable-specific prediction models as the basis for assessing future activity. We evaluated the capability in the context of detecting attacks, demonstrating that our approach can reliably detect direct attacks on the process control. As well, we further explored its potential to identify more sophisticated indirect attacks on field device measurements.

### Inferring Asset Aliases

From a network vantage point, it can prove difficult to track some types of assets which, due to their nature, exhibit "aliasing", i.e., either multiple assets appearing to the network as a single entity (such as distinct hosts behind a NAT all manifesting the same IP address), or the same asset manifesting different identities at different points in time (such as the same mobile device connecting to different parts of an enterprise network, and via DHCP with different IP addresses). In this ongoing project we have



undertaken the discovery of "trackers", i.e., identifiers in network traffic that reveal distinct identity separately from a source's IP address information.

We endeavor to tackle this problem in a fundamental manner by developing algorithms to comprehensively discover trackers of one particularly common class: recurring bytestrings. We start from network traffic for which we have ground truth regarding the actual identity of potentially aliased devices within it. For each such entity, we conceptually extract all the strings the entity transmits, and then filter out any strings that we observe any other entity also transmitting. We look for *recurring* trackers by filtering out any strings a given entity only sends once, and *demonstrative* trackers by filtering out any strings that an entity sends to a destination with which none of the other entities in our pool communicate. At the end of this process we have those strings that each entity sends to a common destination (one with which other entities also communicate) that are unique to the given entity. These then provide high-quality candidates for an analyst to inspect, to assess the context within which the string was transmitted, in order to determine whether the candidate indeed structurally corresponds to a viable tracker.

At present we have the basic algorithms developed and performant enough to process traffic captured for more than a hundred distinct entities over a 10-day period. We are now undertaking the manual analysis stage regarding the context in which the potential trackers appear, and working on further scaling up the algorithms to work on much larger sets of entities, since the approach gains power with the number of entities to which we can apply the analysis.

## THRUST 2: Obtaining dependencies between missions and assets

Cyber-missions are sets of tasks that must be performed in a specified order and within a specified time frame to enable and to support operational missions. These tasks can assume a hierarchical relationship. That is, there are high-level tasks that are composed of a series of more specific, low-level tasks, which could possibly be further broken down into even lower-level ones.

We continued tuning the Sarsia system to better detect dependencies. We also refined and tuned Rippler, which is a complementary system that extracts dependencies based on timing patterns. We continued our research on incorporating external information to augment local network perspectives, and have nearly completed our system infrastructure for very large scale archiving of enterprise activity for later use, such as for "what-if" analysis. We also undertook the development of a novel framework for computing statistical properties of network traffic using distributed sensing.

### Sarsia: a Passive Dependency Detection System

Sarsia is a passive dependency detection system that has several advantages: it is application independent, it is network based, and it does not create any traffic. Because of these properties, Sarsia can be easily deployed in unknown environments. The main disadvantage of a passive approach, like Sarsia, is the fact that it cannot distinguish causal relations from correlated activities. During this reporting period we applied Sarsia to some preliminary data that we received from ABB Control Systems in Norway. We also continued to tune the system to better detect dependencies.

## Rippler: an Active Dependency Detection System

Although passive dependency detection is easily deployable, it also has its own disadvantages. First, passive approaches detect correlated activity and not necessarily causal relationships. Secondly, passive approaches do not provide the direction of dependency. That is, while the output of the passive dependency provides pairs of services as depending on each other, it cannot identify which service depends on the other.

To address this issue in Sarsia, we introduced active approaches. However, to assure that the active system was easily deployable, we needed the active approach to be application-independent and network-based. To achieve this, we implemented an active dependency detection system called Rippler. This system injects small delays in the beginning of network connections to the under-study services. These delays get propagated to the depending and dependent services. We used statistical (hypothesis testing) analysis to detect these propagated delays and to analytically provide appropriate tuning of the system thresholds for gaining desired false positive and false negative rates. We also analytically provide formulas to give insight into the effects of network noise and jitter, overloaded services, cached services, and popular services in the precision (false positive and false negative rates) of the detection mechanism. Furthermore, we showed that no matter how large the noise in the network, arbitrarily small false positive and false negative rates can be achieved by running the experiment for a long enough time. We also analytically calculated the experiment length needed to achieve given false positive and false negative rates and the noise parameter.

We confirmed the correctness of these analyses by simulating networks, containing service dependencies, with different amounts of noise. We installed Rippler in a university lab and injected delays for 5 months. By analyzing the gathered data (in NetFlow format), we detected 54 dependencies. We compared the results with three previous passive dependency detection systems, NSDMiner, Orion, and Sherlock. We showed that Rippler, as expected from an active tool, outperforms these passive approaches. To the best of our knowledge, Rippler is the first application-independent, network-based dependency detection system. During this reporting period we tuned the system and published a paper on Rippler at INFOCOM 2014.

## Adding a Global Vantage Point to the Local Mission

A basic question regarding effective monitoring of local missions concerns the power of augmenting the purely local network perspective by incorporating external information. We continued our case study of such augmentation to understand the degree to which having a broader perspective on relevant activity outside of the local realm can benefit the task of assessing the semantics of activity encountered during a mission. Examining the Web's SSL landscape, we studied the dynamics of global CA trust relationships over time, with the goal to identify patterns which, when seen locally, indicate malicious activity like man-in-the-middle attacks abetted by certificate substitutions. To that end, we set out to understand to which degree *benign* changes to the certificate ecosystem share structural properties with attacks, based on a large-scale data set of more than 17 billion SSL sessions. We found that common wisdom falls short in assessing the maliciousness of an unknown certificate, since numerous artifacts that would seemingly indicate malicious intent in fact routinely occur in benign contexts as well. We also examined what impact our observations have on proposals aiming to improve the security of the SSL ecosystem, finding that in particular Google's Certificate Transparency--while clearly representing a immense step forward for protecting today's fragile trust relationships---fails to address some of the issues we identify.

## Distributed Summary Statistics for Network Awareness

From a mission perspective, monitoring the network of a site often requires correlation of distributed measurements to assemble a comprehensive picture of its activity. We developed a novel framework for calculating a wide array of summary statistics in real-time, independent of the underlying data, and aggregated from independent monitoring points. We focused on providing a transparent, extensible, intuitive interface and implemented our design on top of an open-source network monitoring system. We demonstrated a set of example applications for profiling and statistical anomaly detection that would traditionally require significant effort and different tools to compute. We have released our implementation under BSD license and collected experiences from real-world deployments in large-scale network environments.

## Highly Scalable Archiving of Enterprise Activity

Our "VAST" system (Visibility Across Space and Time) aims to enable very large scale archiving of activity seen within an enterprise (both network and end-host) for purposes of detecting incipient problems (failures, attacks) and facilitating post-facto analysis (forensics, what-if scenarios). The associated analysis tasks require a system that can support interactive lookup of large data corpora spanning extensive periods of time.

During the reporting period, we undertook extensive efforts to bring the system's architecture and implementation to a stable, operationally viable point. One particular focus concerned designing and implementing VAST's query framework, to complement the ingestion and indexing framework already in place from our previous efforts. Overall, VAST is now so developed that an invited presentation on it at the Lawrence Berkeley National Laboratory led the LBNL operational security team to request working with us to deploy VAST for operational use during the coming months. Just after the reporting period, we also presented the system at BroCon, attended by 150 users of our group's open-source "Bro" network monitoring framework, and as a live demo at SIGCOMM'14. We now have our first external users experimenting with the system, reporting back their experience and filing tracker tickets with feature requests.

Our architectural work illuminated several significant issues for such systems. First, message passing systems suffer from the same difficulties with "buffer bloat" as do networks, even when communication occurs within the same process. Initially, VAST buffered and compressed arriving events into data segments of up to 128 MB before relaying them to the indexing engine, which caused high variance and jitter in the pipeline. By switching to a much smaller batch size and deferring the segmentation to the point when data gets written to disk, we achieved a much more uniform load factor of the system under stress. Second, we identified the importance of employing state machines in such systems to handle complex asynchronous messaging patterns. The key building block of VAST, the C++ Actor Framework (CAF), has first-class support for expressing state machines, and its abstractions enabled us to eradicate several difficult-to-reproduce race conditions that arose due to the inherent non-determinism in such highly concurrent but asynchronous systems. Third, we distilled workflows for efficient concurrent many-core programming that resemble GPU-style programming. Using CAF's copy-on-write messages, we build indexes in parallel by sending the same data to multiple, different actors, each of which acts on a slice of the data.

VAST now comprises 37,000+ lines of extensively unit-tested C++14 code. With low-level performance optimization work still pending, we already can import data at 100,000 events/sec for a single system, and achieve sub-second response times when querying data.

## THRUST 3: Obtaining an accurate view of the impact of cyber-attacks

Based on the available information on the cyber-assets that are required by each mission, our framework will draw meaningful conclusions about the current status of the missions being executed and the threats that different attacks pose. This analysis should generate a number of possible courses of action (COAs), highlighting cost-benefit tradeoffs.

Our novel framework for assessing the impact of cyber-attacks handles triaging, and has been tested on a realistic, phase-based cyber-warfare exercise. We developed a system, dubbed Nazca, for picturing the download-behavior of executables within large-scale networks. We presented another system, Hulk, to detect malicious behaviors within browser extensions. Finally, we interviewed and began testing real-world industrial control networks to glean more information on realistic attack-preparedness and awareness.

### Modeling attacks through controlled cyber-warfare exercises

In traditional live competitions, the participants are given virtualized hosts that run services with vulnerabilities. The participants then identify the vulnerabilities, patch their version of the services, and use the knowledge about the vulnerabilities in order to develop an exploit, which is then run against the services of the other participants. In this case, the attack code is controlled by the attacker(s) who developed it. Therefore, it is very difficult to replicate the attack at a later time, in order to recreate a cyber-attack scenario.

In the previous reporting period, we developed a novel framework, called Cyber-Mission-Range (CyMiR). The framework decouples the various components involved in a cyber-attack, such as the assets, the users, the attackers, the exploit code, etc. By doing this, CyMiR supports the *ad hoc* selection of assets, users' behaviors, and types of attacks, which are then used to synthesize a specific attack against a cyber-mission.

This year we extended the CyMiR system to support cyber-triaging, which is the process of assessing cyber vulnerabilities and understanding which are the most critical flaws that need to be addressed in order to protect the network infrastructure.

The system uses state exploration and scenario-based evaluation in order to determine which asset, in a possible scheduling of attacks and services accesses, represents a substantial "security bottleneck" (meaning that one successful attack might affect multiple missions and activities).

In addition, we explored how the CyMiR can be leveraged to not only represent missions, but also explore mission phases and how different mission phases could be targeted by different attackers with different skill-sets.

To test our approach, we created a security exercise that featured a number of phases simulating the steps necessary to acquire the assets to perform a nuclear attack. At each step, each team carries out one or more tasks while being under attack by the other teams. Once a team completes a step, the team is moved to the next step. Teams in a step could only attack teams in the same step. This created a situation in which, as the experiment progressed, teams would be confined in scenarios where they would compete with similarly-skilled teams. This setting models the fact that in cyber-security it is often the case that assets at different levels of criticality are attacked by actors with very different resources and skills.

The competition was a success, with more than a hundred teams participating (and more than a thousand students attacking and defending their assets). The dataset collected will be the basis for a layered analysis to cyber-attacks and defense, as well as the analysis of the effectiveness of cyber-triaging.

## Analyzing the Downloads of Executable Programs in Large-Scale Networks

In this work, we studied how clients in real-world networks download executable programs. Users in large networks frequently download and install new programs as well as updates and patches to existing software. They install applications to work on new tasks, to improve their productivity, and to access content on the Internet. Unfortunately, not all programs that are installed are also desirable. Drive-by download attacks exploit browsers and abuse their functionality to download and install malware programs. Malicious web pages use social engineering to trick a user to access and execute a Trojan horse that compromises the victim's machine.

From a network point of view, connections that fetch a malware program are hardly suspicious, and they look essentially identical to legitimate requests performed by users who download benign programs. However, the situation changes significantly when "zooming out" and leaving the myopic view of individual malware downloads. Instead, when considering many malware downloads together -- performed by different hosts, but related to a single campaign -- a malware distribution infrastructure becomes visible. In some sense, this malware distribution infrastructure acts like a content distribution network. However, there are also differences, and these difference can be leveraged to identify cases where malicious content (malware programs) are distributed.

To detect connections (more specifically, HTTP requests) that are used to download malware binaries, we developed a system called Nazca. Similar to the drawings in the Nazca desert, the malware downloads (and the supporting distribution infrastructure) become more apparent when observing a larger part of the picture/network. Our decision to focus on HTTP (web) connections is driven by the observation that an overwhelming majority of drive-by exploits and social engineering attacks use the web to download malware binaries.

Our system monitors web traffic between a set of hosts (typically, machines in the protected network) and the Internet. The goal is to identify those connections that are related to malware downloads. To this end, Nazca operates in three steps: In the first step, the system identifies web (HTTP) requests and extracts metadata for subsequent analysis. This metadata includes information about the connection endpoints, the URIs, and whether an executable program was downloaded.

In the second step, our system identifies suspicious web connections. A suspicious connection is a web request that downloads an executable file, and there are certain properties associated with the connection that make it appear different from legitimate downloads. These properties are designed to capture techniques employed by malware authors to hide their operations from traditional defense systems. Such

evasive techniques include domain fluxing, malware repackaging, and the use of malware droppers for multi-step installations. An interesting, and favorable, trait of our approach is that it complements well existing detection mechanisms. That is, when malware authors employ techniques to evade traditional approaches (such as malware signatures or IP/domain reputation systems), their downloads are easier to recognize for Nazca.

In the third step, Nazca aggregates the previously-identified suspicious connections (candidates). The goal is to find related, malicious activity. The purpose is to reduce potential false positives, and to focus the attention to the most significant infection events. Moreover, as a byproduct of this step, we build a graph of the malicious activities that our system detected. This enables the reconstruction and observation of entire malware distribution networks.

To make detection decisions, our system neither looks at properties of the downloaded programs nor at the reputation of the host that serves the program. Thus, Nazca does not suffer from coverage gaps in reputation databases (blacklists), and it is not susceptible to code obfuscation. To demonstrate the effectiveness of Nazca, we evaluated it on a two-day and one-week traffic dataset from a large Internet Service Provider. Our results show that the system is effective in detecting malicious web downloads in large-scale, real-world environments.

## Detecting Browser Subversion Threats

Recently a particularly serious vector for compromising assets has emerged in the form of malicious browser extensions. Just as app stores facilitate vast additional functionality for mobile devices, markets offering 10-100s of thousands of browser extension markets provide a means by which users looking to enhance their browsing experience can introduce into enterprises a wide range of altered functionality. Unlike traditional downloads, however, this occurs in a manner to which much security monitoring remains blind, because extensions execute wholly within the browser's internal environment. Attackers do not need to introduce new executables into end systems to take effective control over them in terms of hijacking a user's browsing - they just lure users with enticing browser extensions, or, more insidiously, they buy (or otherwise take over) the development rights to an existing, widely used extension and use that to push out "updates" that introduce arbitrary malicious functionality. Some extensions have more than 10 million users.

To analyze this threat and develop technology to counter it, we developed "Hulk", a dynamic analysis system that detects malicious behavior in browser extensions by monitoring their execution and corresponding network activity. Hulk elicits malicious behavior in extensions in two ways. First, Hulk leverages HoneyPages, which are dynamic pages that adapt to an extension's expectations in web page structure and content. Second, Hulk employs a fuzzer to drive the numerous event handlers upon which modern extensions heavily rely. We analyzed 48,000 extensions from the Chrome Web store and identified 130 as clearly malicious - including one with more than 5 million users - and more than 4,700 as potentially malicious.

## Interviewing Large-Scale Critical Industrial Control System Operators

To better assure that our cyber-awareness framework will give enterprise users the environment that they need, we interviewed system operators of critical large-scale industrial control networks that need to be available 24/7. Our intent was to assess how prepared the operators were for dealing with cyber attacks.

One of the biggest problems we encountered was finding entities (government or commercial) that were willing to let us interview their operators. After some initial disappointments, we found some inroads in the electrical power industry.

Targeted cyber attacks are on the rise, and the power industry is an attractive target. The two most likely goals of targeted attacks are espionage and causing physical damage. In the case of the power industry, the worst possible consequences are severe: large areas, including critical societal infrastructures, can suffer from power outages. During this review period we ran a study to measure the preparedness of the power industry against targeted attacks. We interviewed six power distribution system operators (DSOs) to assess the level of cyber situation awareness among the DSOs and to evaluate the efficiency and effectiveness of their current deployed systems and practices for detecting and responding to targeted attacks. Our findings indicate that the power industry is very well prepared for traditional threats, such as physical attacks. However, cyber attacks, and especially sophisticated targeted attacks, where social engineering is one of the strategies used, have not been addressed appropriately.

Our plan is to improve the design of our cyber-awareness framework based on the information that we gathered in the tests performed during this reporting period. Also, by reviewing previous targeted attacks and learning from them, we hope to provide industry with guidelines for improving their current situation awareness and defense capabilities.

## THRUST 4: Obtaining actionable cyber-attack forecasts

The cyber-awareness domain possesses significant challenges to mainstream game theory, which can only be overcome through fundamental research in this area. While it will often be possible to determine unequivocally whether or not a particular task was successfully accomplished, this may not always be the case. Conversely, the adversary may also not be able to determine accurately the current state of the mission. In fact, the mission's success may to a great extent rely on this. Partial information games, in which one or both players are not fully aware of the current "state" of the game, are especially challenging and, while optimal solutions may be impractical, it is possible to construct solutions that are suitable for cyber-attack forecasting. Another important challenge is that the complete set of adversary actions (i.e., attacks) will typically not be known a priori. Additionally, one may be faced with large uncertainty regarding the information available to the adversaries and even the ultimate intent behind their actions.

We extended and refined our approaches to asymmetric games, addressed the solution of large-scale games using randomized methods, proposed some optimization-based algorithms for attack prediction, developed some measurement models for estimating the status of cyber assets, and developed some new approaches to the consensus problem when under attack.

## Asymmetric Information Games and detection

At Georgia Institute of Technology, we have continued our work on asymmetric information games. The setup consists of two actors (e.g., defender vs attacker), in which one actor has more information than the other. The informed actor is faced with the dilemma of revelation vs exploitation, i.e., exploiting superior information can result in revealing this information to the uninformed actor. This problem falls under the game theoretic setting of "repeated games with partial information". These results complement work by

UCSB team members in that it provides “worst case” guarantees against an arbitrary opponent (instead of an opponent with limited computational capability).

We have expanded on this work significantly to accommodate more general settings. These include time varying environments and stochastically evolving environments, where the state evolution depends only on the informed player’s actions. As before, we develop methods based on reparameterization and linear programming that enable low complexity solutions for the informed player. We illustrate the approach on a traveling inspector problem.

We also are continuing our work on model reduction for hidden Markov models. Recent work investigates lower bounds on achievable reduction and reduction of models with control inputs (i.e., decision problems).

Finally, we have initiated an investigation into repeated games for LTE security. In this setup, there is an attacker in the form of a smart jammer who wishes to either bring down the network or unfairly exploit network resources. We develop a game theoretic formulation to investigate the ability of a system planner to thwart such attacks. In particular, we show the advantage of a repeated game formulation versus a one shot formulation.

## Large-Scale Games

The matrix games that arise in cyber-attack forecasts are exceedingly large because of the exponential size of the decision spaces both for defenders and for attackers. We addressed the solution of such large games using randomized methods. We developed two algorithms:

1. We formalized a procedure, termed as the *sampled security policy (SSP) algorithm*, by which a player can compute, with a high confidence, security policies against an adversary using randomized methods to explore the possible outcomes of the game. The SSP algorithm essentially consists of solving a stochastically sampled subgame that is much smaller than the original game.
2. We also proposed a randomized algorithm, termed as the *sampled security value (SSV) algorithm*, which computes a high-confidence security-level (i.e., worst-case outcome) for a given policy, which may or may not have been obtained using the SSP algorithm.

For both the SSP and the SSV algorithms we provided results to determine how many samples are needed to guarantee a desired level of confidence. We started by providing results when the two players sample policies with the same distributions, and subsequently extended these results to the case of mismatched distributions. These theoretical results were validated into a family of security games based on the 2010 international Capture the Flag Competition (iCTF).

## Optimization-based Attack Forecasting

Mission-centric cyber-security analysts require a complete overview and understanding of the state of a mission and any potential threats to their completion. To facilitate this, we proposed an optimization-based algorithms that can be used to predict in real-time how an attacker may try to compromise a cyber-mission with a limited amount of resources, based on a model that takes into account potential damage to the mission and probabilistic uncertainty. Two different optimization schemes were considered: one where all the mission data is known a priori to the attacker, and another where system identification



and a moving horizon optimization is used to produce the estimates based on historical data. Our schemes were compared with real attacks carried out by human players in the 2011 international Capture The Flag (iCTF) hacking competition. Optimization schemes also allow a defender to go through the missions, and *automatically* protect the services using shields against an “optimal” simulated attacker. We provide various insights to the defender so that the missions can be prevented from being compromised. These results have been incorporated into the visualization tools and human-in-the-loop experiments that arose from Thrust V.

## Estimation Under Cyber-attacks

Cyber-physical missions typically rely on a multitude of sensors that may be compromised by an adversary. Such services may provide information about the status of cyber assets (e.g., whether or not a server is operational) or about physical entities (e.g., whether or not troops have been detected in a given region). A successful attack on such services may significantly hamper mission success. We addressed the problem of estimating the true status of a service based on a multitude of sensors that may have been tempered by an opponent. We formulate this problem mathematically as the estimation of a binary random variable  $\theta$  based on  $m$  noisy measurements that can be manipulated by an attacker. The problem is formulated as a minimax optimization, where one seeks to construct an optimal detector that minimizes the “worst-case” probability of error against all possible manipulations by the attacker. A significant novelty of our approach with respect to classic problems of Byzantine faults is that we do not assume perfect sensors, i.e., even the sensors that have not been manipulated can produce incorrect results, which is common in the cyber-security domain. The problems considered cover several options regarding the type of measurements considered and the amount of information available to the attacker:

1. The measurement model was extremely general, simply assuming that each of the  $m$  measurements has a known distribution conditioned to the true value of the variable  $\theta$  that one wants to estimate. This can capture binary sensors that would report the true value with high probability, but occasionally report erroneous measurements, as well as sensors that would produce continuous measurements obtained by adding random noise to the true value. Our results have been specialized to binary sensors that seem most relevant to this MURI project.
2. In some cases, we made the worst-case assumption that the attacker has full information about the true value of  $\theta$  and all the  $m$  measurements. In other instances, we assumed that the attacker knows the true value of  $\theta$  but does not know the values reported by the sensors that have not been compromised. These different assumptions on the information structure of the game resulted in different optimal estimators.

## Distributed Consensus Under Cyber-attack

Numerous cyber-missions rely on distributed algorithms to solve the *consensus problem*, i.e., achieving agreement among a number of processors for a single value of data. It is generally assumed that some processes may fail or communication may be unreliable so the algorithm must be fault tolerant. We have developed consensus algorithms that aim at detecting misbehaving agents and minimizing their influence on the final outcome:

1. In two papers we addressed the problem of a consensus system in the presence of Byzantine faults seen as an attacker injecting a perturbation in the state of the nodes. We used Set-Valued Observers to detect if the state observations are compatible with the system dynamics. The method was extended to

the stochastic case by introducing a strategy to construct a set that is guaranteed to contain all possible states with, at least, a pre-specified desired probability. The proposed algorithms are stable in the sense that it requires a finite number of vertices to represent polytopic sets while also enabling the a priori computation of the largest magnitude of a disturbance that an attacker can inject without being detected. Moreover, the algorithm proposed in our paper, "Finite-time Average Consensus in a Byzantine Environment Using Set-Valued Observers", achieves convergence in finite time.

2. We presented a game theory-based consensus problem for leaderless multi-agent systems attempting to agree upon a time-varying quantity and its 1<sup>st</sup> derivative in the presence of adversarial inputs that corrupt the measurements. The problem is addressed under a distributed decision making framework that is robust to possible cyber-attacks. For each agent, we derive three tuning laws: one associated with the cost, one associated with the controller, and one with the adversarial input. These ideas were subsequently pursued in a number of our publications to construct distributed decision making algorithms that use learning to overcome system uncertainty.

## THRUST 5: Obtaining a semantically-rich, easy-to-grasp view of the cybermission status

Successful information transfer to a decision-maker and the right analysis tools given a particular user-context are crucial components of superior cybersecurity situational awareness. Collecting, extracting, mapping, filtering, modeling, and predicting security data are all important steps for creating a safe, secure, and resilient cyberinfrastructure to support and protect important missions. But without the consideration of context-specific knowledge dissemination theory and mechanisms, the decision maker will not be able to arrive at the correct analysis, understanding, and dynamic explanation of such data, or unveil the options and consequences for the decisions ahead, and do all of this in a timely fashion before the window of opportunity for appropriate counter-measures closes. An effective approach to true cybersecurity-awareness requires a multi-disciplinary collaboration among security researchers and practitioners, game theoreticians, and visualization and user interface experts. There is a need and distinct opportunity for scientific advances in interactive visualization methodology that scales appropriately with regard to changes in display and interaction capabilities, and mission context.

Over the past review period, we have made significant progress in three overall areas: 1) We updated our surround-view immersive situation room and visualization chamber, the UCSB Allosphere, to full-surround projection of 2D, 3D, and multimodal information displays. 2) We have created flexible tools for visual and multimodal cybersecurity data analysis. 3) We have been working towards automatic generation of user interfaces for different display platforms by conceptualizing a universal syntax for data and user interfaces, providing structured editing tools for data and UIs, and targeting transparent "deconstructible" user interfaces and visualizations.

### Cybersecurity Situation Room

For our situation room, the UCSB Allosphere, our hardware and software infrastructure now supports full-surround presentation of 2D or 3D information displays, including multi-modal elements. We support visual, audio and interactive representation, transformation and generation across a diverse set of content areas. As reported previously, we have also evaluated full-surround stereoscopic visual design as well as 3D spatial audio to increase immersion in the instrument. Visual calibration has been a key component

over this review period, and we have achieved a seamless view across the multiple projectors lighting the sphere surface. Multi-user interaction using a variety of devices has been another active area of research. We believe that all these affordances facilitate immersive, multi-user participation for cyber situational awareness.

## Flexible Interactive Multimodal Information Visualization

We have integrated tools that enable users to design information visualizations for complex datasets interactively, without the need for textual programming. Our system, iVisDesigner, achieves high interactive expressiveness through conceptual modularity, covering a broad information visualization design space. iVisDesigner supports the interactive design of interactive visualizations, such as provisioning for responsive graph layouts and different types of brushing and linking interactions.

Our system was designed to be a flexible tool for interactively creating information visualizations, inspired by interactive vector-based drawing tools and established information visualization principles. The system allows users to freely place graphical elements, and links between them, on a large surround canvas. We chose a declarative approach to avoid reliance on familiarity with programming and for keeping the usage simple and straightforward. Our unified editing interface allows users to create and edit graphical, guide, and generator objects, enabling the interactive design of complex visualizations.

Our design for multimodal integration comprises a consistent notation across modalities in addition to high-level abstractions affording intuitive declarations of multimodal mappings, unified timing constructs, and rapid, iterative re-inocations of constructors while preserving the state of audio and visual graphs. This work builds on our design of a Javascript-based programming language, Gibber, which provides several programming abstractions that simplify exploratory programming practice and multimodal authoring.

## Towards Self-Descriptive Data Formats and Platform-Scalable User Interfaces

There is a classic design tension between user-friendly user interface design and expert-friendly user interface design. There is also a classic design tension between binary data format design and printable data format design. Our work on automatic generation of user interfaces attempts to address both sets of design tensions as having the same cause and solution. We observe an opportunity to redefine the baseline for “human-readable” data formats by pairing a universal binary syntax with a universal structured editor and explore the rippling implications that it could have on human-computer interaction and user interfaces on various platforms. Once all data has inherent metadata capacity, markup languages can be mixed in to augment any data at any time with high-level types and presentation descriptions. Object-oriented agents can then be automatically attached to individual data portions to provide various layouts, abstractions, visualizations, or interactive APIs that uphold invariants across operations. Agent implementations live on the client and take care of mapping to platform specific modalities. This paradigm is about feeding a multitude of micro applications and user interfaces to data as opposed to the other way around. These data-driven layers can be combined and layered recursively, maintaining the provenance of each model and view’s inputs. One can imagine incrementally approximating a full-fledged traditional graphical user interface, which can then, by construction, be peeled back apart on the fly for any number of reasons.

To a developer, there is value in prototyping and testing each layer in an application's architecture incrementally. Thus, there is an incentive to approach development as structural data modelling with adaptable views and APIs for each milestone, starting very general and finishing very specialized. This gives applications, such as our cyber situational awareness interfaces, a spreadsheet-like quality, suitable for end-user programming.

## Synergistic Efforts

Finally, we continued our work on cognitive models of Trust and Situation Awareness in synergy with the ARL-sponsored Network Science Collaborative Technology Alliance (NS CTA). Using cognitive and user modeling, as well as game-theoretical strategy analysis, we researched the interdependencies of human trust behavior, situation awareness, and performance within a realistic multi-genre network involving humans, information sources, and intelligent agents. Through controlled experimentation involving social dilemmas and military context scenarios, we studied the emergence, influence factors, and potential control of trust and related cognitive phenomena in composite networks, mediated by user interface technologies.